

Automated Radiological Impression Generation for Plain Chest X-rays with End to End Deep Learning

Shuai Zhang^{1,2†}, Xiaoyan Xin^{1†}, Jingtao Shen³, Yachong Guo^{2,4}, Yang Wang¹, Xianfeng Yang¹, Jun Wang^{2,4,5}, Jian Zhang^{2,4,5*}, Bing Zhang^{1*}

1. Department of Radiology, Nanjing Drum Tower Hospital, the Affiliated Hospital of Nanjing University Medical School; Nanjing 210008, P.R. China

2. School of Physics, Collaborative Innovation Center of Advanced Microstructures, Nanjing University, Nanjing, China

3. Department of Nuclear Medicine, Nanjing Drum Tower Hospital, the Affiliated Hospital of Nanjing University Medical School; Nanjing 210008, P.R. China

4. Institute for Brain Sciences, Nanjing University, Nanjing, China

5. State Key Laboratory for Novel Software Technology, Nanjing University, Nanjing, China

† Shuai Zhang, Xiaoyan Xin contributed equally to this work and considered co-first authors.

*Correspondence authors:

Jian Zhang, jzhang@nju.edu.cn. Bing Zhang, zhangbing_nanjing@nju.edu.cn;

Abstract

The chest X-Ray (CXR) is the one of the most common clinical exam used to diagnose thoracic diseases and abnormalities. The volume of CXR scans generated daily in hospitals is huge. Therefore, an automated diagnosis system that is able to save the effort of doctors is of great value. At present, the applications of artificial intelligence in CXR diagnosis usually use pattern recognition to classify the scans. However, such methods rely on labeled databases. They are costly and usually have a high error rate. In this work, we built a database containing more than 12,000 CXR scans and radiological reports, and developed a model based on deep convolutional neural network and recurrent network with attention mechanism. The model learns features from the CXR scans and the associated raw radiological reports directly; no additional labeling required. The model provides automated recognition of given scans and generation of impression. The quality of the generated impression was evaluated with both the CIDEr scores and by radiologists as well. The CIDEr scores were found to be around 5.8 on average for the testing dataset. Further blind evaluation suggested a comparable performance against radiologists.

Keywords

Chest X-ray, Machine learning, Deep learning, Neural networks, Attention mechanism

Introduction

The chest X-Ray (CXR) is one of the most common diagnostic techniques for respiratory system. It is quick and inexpensive, and yields low radiation. The volume of daily CXR scans in hospitals is huge and their examination and interpretation consume lots of time and effort of radiologists. Therefore, it is desirable to develop an automated system that is able to examine and interpret CXR radiographs automatically. Moreover, an automated system may help reduce inter-observer variations due to the factors including individual experience, quality of the radiograph, time and personality type [1]. The adoption of an automated system will lead to a more standardized terminology and treatment, and benefit the collaborations between different parties. The system may further evoke new applications such as remote diagnosis, self-service diagnosis and so on.

Previously, many works focused on automated classification of the CXR scans. These works are usually based on variants of Convolutional Neural Networks (CNN) and supervised learning [2-9]. However, at least three problems hinder the practical applications of automated systems in hospital with these methods. First, the label of the chest film is usually extracted from the report, and its accuracy is not guaranteed. Second, the sensitivity and false positive rates of these classification approaches had been saturated. After all, many diseases can't be distinguished by experts if they only look at chest films. Third, the decision strategy underlying these systems has not been well understood, making it difficult to track the errors and gain the trust of doctors and patients.

In this work, we developed a model based on deep convolutional neural networks and recurrent neural networks with attention mechanism that learns from CXR images and the raw radiological reports simultaneously. The deep neural networks have shown great potential in characterizing and classifying complex data in a broad range of fields [10, 11, 12]. After training with our database, the neural network is able to automatically generate radiological impression for given scans. Our work has the following novel features. First, the training of our model is only weakly supervised; the model directly learns from the image and the raw radiological reports stored in the hospital database; no further classification and labeling of the images by human is required. That is, in contrast to most machine learning models. This feature greatly facilitates the acquisition of data and training of large-scale models. Second, instead of a simple classification of the case into one or several disease categories, the output of our model is a descriptive report regarding different conditions of the chest; the output is directly readable for patients. Third, the implementation of the attention mechanism adds another level of understanding of how the model works, facilitating debugging and optimizing of the model. In the following sections, we describe the model architecture, training and testing procedures, and the performance evaluated with the CIDEr scores and by human radiologists.

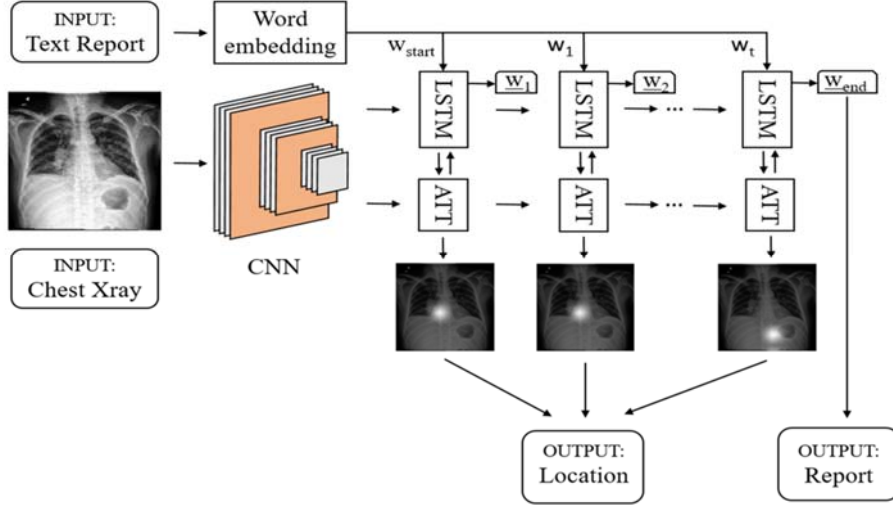
Methods

The architecture of the network

Figure 1 shows the overall architecture of the neural network. During training, the neural network reads in both CXR images and the raw radiological reports, and outputs human readable texts. The output is

then compared with the ground truth to calculate the loss function, which is minimized with the gradient back-propagation technique. After training, the model is able to automatically generate impression for given CXR images. The design of the model architecture was inspired by the pioneer work of Xu et al. [13], where they developed an RNN to generate captions for daily images, such as those in Flickr and MS COCO databases. The model is also similar to those by Zhang et al. [14] and Wang et al. [15] in terms of the purpose of automatically generating medical reports. However, the architecture of our model was redesigned to better fit the organization of our database.

Figure 1: The architecture of the whole network



The model contains a 121-layer Densely Connected Convolutional Network (DenseNet) [16], which is used as a visual information encoder to extract features from the input images. The encoder is composed of four blocks; each block contains several convolutional layers, each takes all preceding feature-maps as inputs. The blocks are connected by transition layers. According to Huang et al., DenseNets alleviate the vanishing gradient problem, strengthen feature propagation, encourage feature reuse, and substantially reduce the number of parameters [16]. Compared with many other CNNs, they converge faster and are appropriate for smaller datasets. Therefore, DenseNets are suitable for medical images. The output of the last layer of the DenseNet block is fed to the Long Short-Term Memory (LSTM) network to generate descriptions for the given CXR image.

The LSTM network [17] is adopted to generate texts for the given CXR image word by word. At each step, it reads the output of the last layer of the DenseNet block and the previous generated word, and outputs the next word. In detail, our LSTM implementation follows that of [18], i.e.,

$$\begin{pmatrix} i_t \\ f_t \\ o_t \\ g_t \end{pmatrix} = \begin{pmatrix} \sigma \\ \sigma \\ \sigma \\ \tanh \end{pmatrix} W_{h,e+c+h} \begin{pmatrix} Ey_{t-1} \\ V_{gav} \\ h_{t-1} \end{pmatrix}, \quad (1)$$

$$c_t = f_t \odot c_{t-1} + i_t \odot g_t, \quad (2)$$

$$h_t = o_t \odot \tanh(c_t), \quad (3)$$

where i_t , f_t , c_t , o_t , h_t are the input, forget, memory, output and hidden gates of the LSTM, respectively. W is the weight matrix, σ is the logistic sigmoid function, V_{gav} is the global average of DenseNet output, y_{t-1} is the previous generated word, and E is an embedding matrix. The symbol \odot denotes element-wise multiplication. The subscript h is the size of the hidden states, e is the vocabulary size, and c is the channel number of the output of DenseNet.

Attention mechanism has been widely adopted in visual image processing since it improves the model performance and adds a level of understanding of how the model works. It mimics the human visual attention mechanism by learning to focus on a certain image region. Specifically, a soft attention mechanism is implemented in the model, which calculates a set of weights conditioning on the image representation and on the hidden state. These weights are multiplied with the output vectors of the DenseNet to get a weighted representation of the image, which is then utilized by the recurrent neural network to generate descriptions. The corresponding equations are

$$\alpha_t = \text{softmax}(W_{p,h}h_t + W_{1,c}V) \quad (4)$$

$$C_t = V\alpha_t \quad (5)$$

$$y_t \sim p_t \propto \exp(W_{e,h}h_t + W_{e,c}C_t) = \text{softmax}(W_{e,h}h_t + W_{e,c}C_t) \quad (6)$$

where α_t is the attention weight, V is the output of DenseNet, C_t is the weighted context vector, p_t is the probability distribution of the words, and y_t is the predicted word sampled from p_t .

The loss function is the cross entropy between the ground truth and the prediction distributions of the texts,

$$\text{loss}(\mathbf{p}, \mathbf{y}) = -\frac{1}{l} \sum_{j=1}^l \log(p_j(y_j)), \quad (7)$$

where p_j is the probability distribution predicted at the j -th step, y_j the index of the j -th word in the ground truth text, and l is the length of the text.

Datasets

All the chest X-ray scans and the associated radiological reports were provided by Nanjing Drum Tower Hospital. In total the dataset contains 12, 219 images and the same number of reports in Chinese language. Among them, 7516 cases were from outpatient department and 5063 cases were from physical examination. According to the corresponding reports, there were 7547 normal cases and 4672 abnormal cases. All reports were investigated by one expert (of attending doctor level or above) and double checked by another expert (of associate chief physician level or above). The dataset was randomly split into three sets, 80% samples for training, 10% for validating, and 10% for testing. Figure 2 shows some examples randomly taken from the dataset.

Figure 2: Four samples chosen from the dataset



No obvious abnormalities in the whole chest.



Increased lung markings and enlarged heart shadow.



Bronchitis, old lesions in the upper part of the lung, right pleural effusion.



Hydropneumothorax on the right side, the lung tissue is compressed by about 40%.

Since there are no blanks between Chinese words, the python module jieba [19] was used for text segmentation. After processing all the radiological reports in the dataset, a vocabulary of size 424 was obtained. The words were represented with one-hot-vectors. The words that appeared less than three times were represented with a special token <nou>. Two other special tokens were <start> and <end>, indicating the beginning and ending of the reports, respectively.

Training Procedures

Transferring learning technique was employed to speed up the convergence of training. Specifically, the 121-layer DenseNet was pre-trained with the ChestX-ray8 dataset released in September 2017 [20] for a classification task in a supervised way. The training procedure was similar to that in [2]. The ChestX-ray8 dataset contains 110k chest x-ray images and 14 types of diseases labels. The trained weights were then transferred to the encoder module of the model. In the training process that followed, the parameters in the first two dense blocks were fixed while that in the others were fine-tuned by the gradient back-propagation algorithm.

During training, the original X-ray images were resized to 256*256 pixels and then processed with histogram equalization to increase the contrast. The size of the hidden unit of LSTM was set to 512 and the embedding size was 256. Adam optimizer was used for the optimization process. The learning rate of the DenseNet was set to 1.0×10^{-4} and that of the LSTM was set to 5.0×10^{-4} .

4. Evaluation metrics

In order to evaluate for image I_i how well a generated sentence c_i matched the consensus of a set of descriptions $S_i = \{s_{i1}, \dots, s_{im}\}$, the Consensus-based Image Description Evaluation score (CIDEr) [21] was used.

To calculate the CIDEr score, the Term Frequency Inverse Document Frequency (TF-IDF) weighting $g_k(s_{ij})$ for each n-gram ω_k in the sentence s_{ij} was first calculated as

$$g_k(s_{ij}) = \frac{h_k(s_{ij})}{\sum_{\omega_l \in \Omega} h_l(s_{ij})} \log \left(\frac{|I|}{\sum_{I_{pq} \in I} \min(1, \sum_q h_k(s_{pq}))} \right), \quad (8)$$

where $h_k(s_{ij})$ is the number of times an n-gram ω_k occurs in the sentence s_{ij} , Ω is the vocabulary of all n-grams and I is the set of all images in the dataset. Then the CIDEr_n score for n-grams of length n was calculated as

$$\text{CIDEr}_n(c_i, S_i) = \frac{1}{m} \sum_j \frac{g^n(c_i) \cdot g^n(s_{ij})}{\|g^n(c_i)\| \|g^n(s_{ij})\|}, \quad (9)$$

where $g^n(s_{ij})$ is a vector formed by $g_k(s_{ij})$ corresponding to all n-grams of length n. $g^n(c_i)$ is similarly defined for the generated sentence c_i .

At last, the CIDEr score is calculated as the average over all n-grams,

$$\text{CIDEr}(c_i, S_i) = \frac{1}{N} \sum_{n=1}^N \text{CIDEr}_n(c_i, S_i). \quad (10)$$

Results

Figure 3(a) shows the losses of the training and validation datasets as a function of the training epoch. It can be seen that the training loss keeps decreasing, while the validation loss saturates at about the 10th epoch, indicating that the generalization ability of the model reaches its maximum. Therefore, the parameters obtained at the 10th epoch are used by the model to generate the results in the following sections.

Figure 3(b) shows the calculated CIDEr values for the testing dataset as a function of epoch. Note that for the testing set, the ground truth sentences were not used when generating the descriptions; they were only used to evaluate the descriptions after their generation. The Beam Search technique [22] was used to generate multiple sentences for a given CXR image, and each sentence was assigned a preference probability. The top three sentences with the highest probabilities were recorded. Their CIDEr values were calculated against the ground truth and the highest one was used to calculate the curve shown in

Fig. 3(b). According to the gure, the average CIDEr value of the testing set increases as a function of the epoch, and saturates around 5.8 at the 10th epoch.

Figure 3: (a) The losses of the training and validation datasets as a function of epoch. (b) The CIDEr values of the testing dataset as a function of epoch.

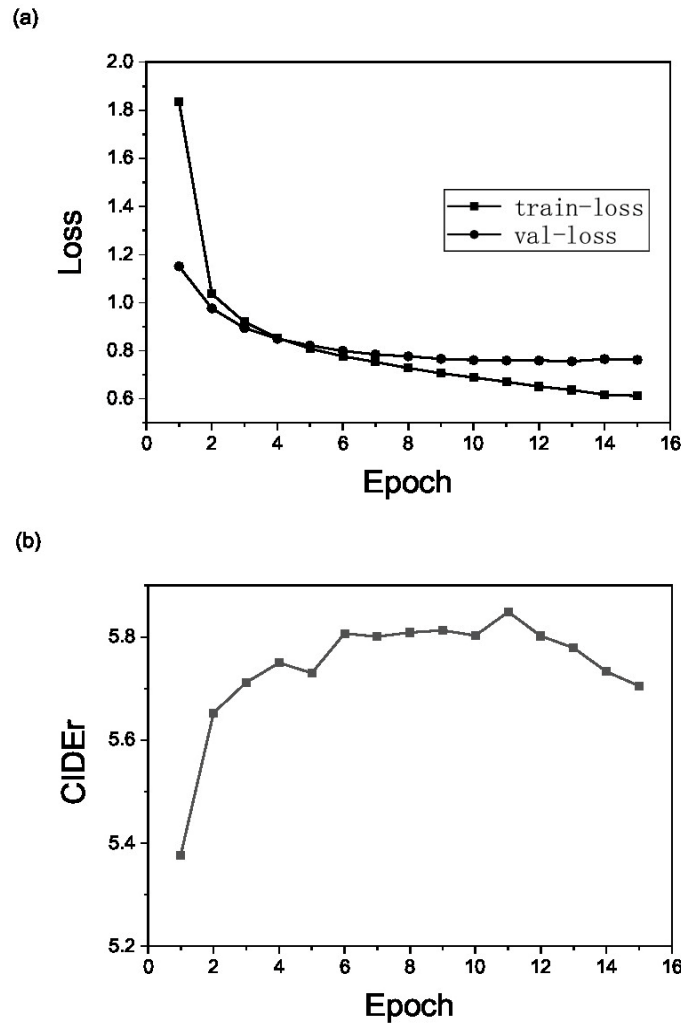


Figure 4 shows several examples of the generated descriptions. For each scan, the top three predictions given by the model are shown, labeled as Pd1, Pd2, and Pd3, respectively, in the decreasing order of the preference probability. More examples are given in the supplemental materials. Fig. 4(a) shows a normal case with increased lung markings in both lungs. The model correctly recognizes the situation and generates descriptions with “increased lung markings in both lungs” in Pd1 and Pd3; while in Pd2, the model says “no obvious abnormalities”. Fig. 4(b) shows the scan of a patient with chronic bronchitis and inflammation, concluded based on the image as well as on the medical history. The model reports “increased lung markings” in Pd1 and Pd2, and directly gives “bronchitis” in Pd3, which is amazing since the model has no information of the medical history. Fig. 4(c) is a case with pleural effusion on the right side. The model identifies the symptom and correctly generates impression for it. For the case in Fig. 4(d), the model correctly recognizes the increased size of the heart and also assumes the image to be a postoperative view. It is not known whether the model draws this conclusion based on the thin and bright strips near the heart region.

Figure 4: Several examples of the chest x-ray scans in the testing dataset and the corresponding descriptions generated automatically by the model

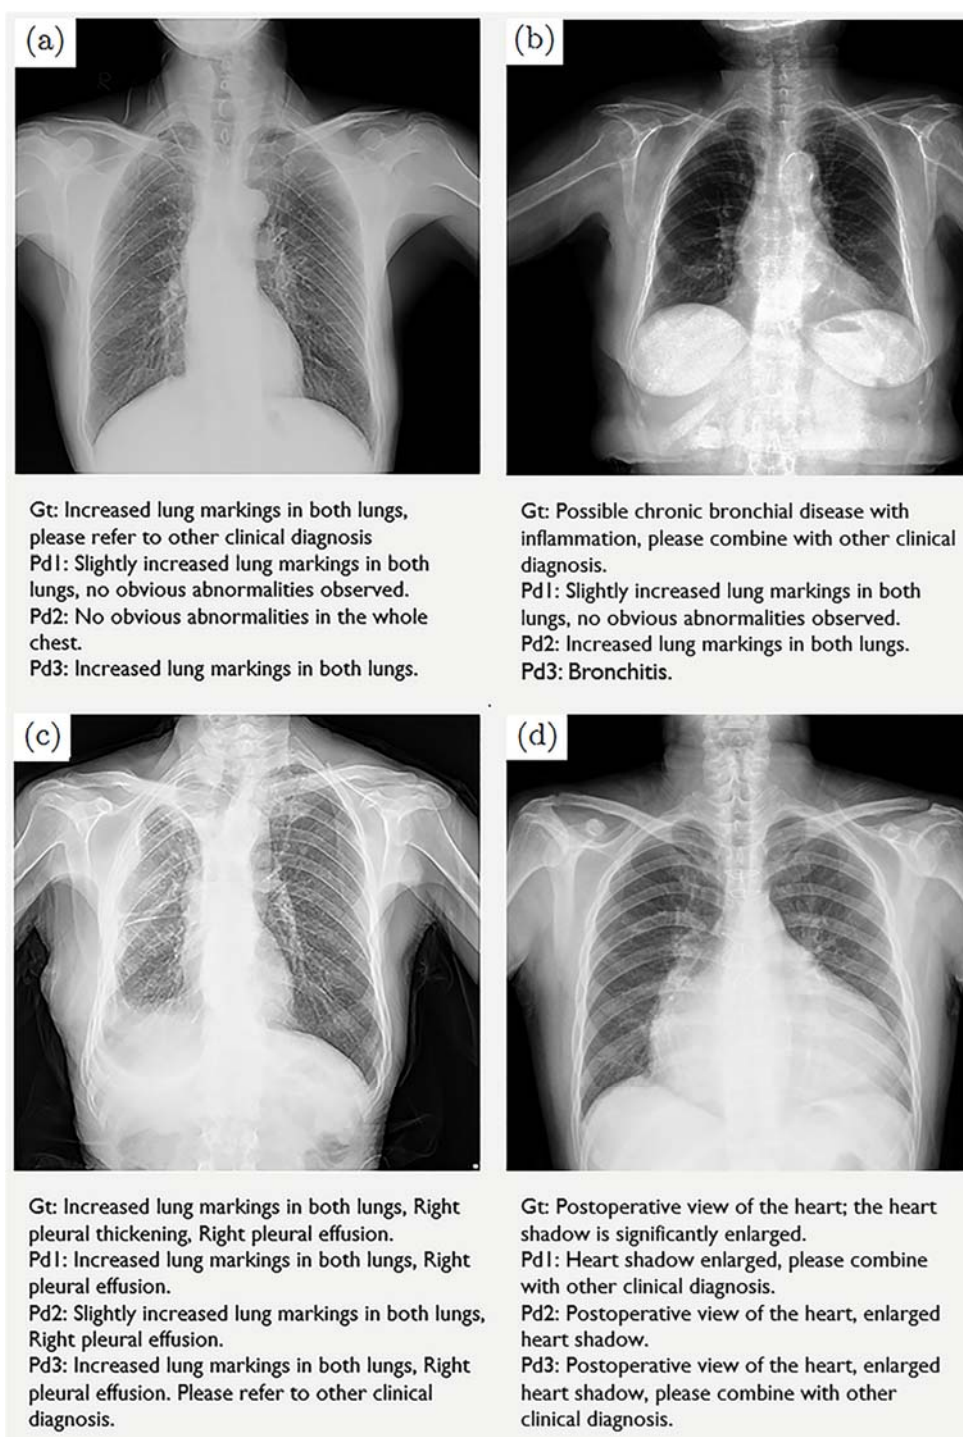



Figure 5 shows the alignment of generated words with the relevant parts of the CXR images. In general, the alignments are consistent with human intuition. The alignments are enabled by the attention mechanism and provide another level of understanding of how the network works. They also facilitate debugging of the results.


Figure 5: Attention over time



The quality of the automatically generated reports were also evaluated by radiologists. The procedure was as follows. 100 CXR scans were randomly extracted from the testing dataset and fed one by one to the model to generate impression. Another 100 CXR scans and the corresponding raw reports, which were written by human radiologists, were randomly extracted from the same dataset. These 200 scans and the associated reports were put together, shuffled and sent to experts for human inspection. Two radiologists (of associate chief physician level) were invited to examine the images and assess the quality of the associated reports, without knowing the origin of the reports - from human or machine. This was to prevent possible bias, either to human or to machine. The radiologists gave scores from 1 to 5 for each report, according to the standard as follows. An report with all conditions identified and accurately described was scored 5; an report with major conditions identified correctly but minor problems outside chest missed was scored 4, e.g., scoliosis, foreign matter in vitro; an report with major conditions identified correctly but minor problems inside the chest missed was scored 3, e.g., old lesions, fibrous stripes, post thoracic surgery, aortic calcification; if major conditions were identified but described inaccurately, a score 2 was given; If major conditions were missed or identified incorrectly, the score would be 1.

Figure 6 shows the score distributions for two groups of reports. It can be seen that for both groups, the majority are scored 5. For the group of reports given by human, 77% are scored 5; while for the reports from the model, 72% are scored 5. At the level of score 4, these two numbers are 9% and 14%, respectively. If we assume the scores of 4 or above are acceptable, then both groups have 86% falling in this range, suggesting that the neural network is able to generate reports with quality comparable with that of radiologists.

Figure 6: Comparison of the quality of two groups of reports



Discussion

In summary, we developed a scheme that is able to automatically generate radiological descriptions for given CXR, based on deep convolutional neural network and recurrent neural network with attention mechanism. We built a database containing more than 12,000 CXR scans and trained the model, and then evaluated the quality of the generated descriptions. The comparison between the generated descriptions and the ground truth gave a CIDEr value of 5.8. We also blended the generated descriptions with that given by radiologists, and invited other radiologists to score them. It was found that among the reports given by radiologists, 77% received the highest score 5; while for the reports generated by the model, 72% were scored 5. For the reports with score 4, the percentages were 9% and 14%, respectively. Therefore, the model is able to generate reports with high quality comparable to that of radiologists, and has the potential to be significantly improved as more training data are available.

The model developed here has some particular features. First, it learns from the raw radiological reports and is able to directly utilize the huge volume of CXR data generated in the hospital; no additional labeling work is required. This feature is particular useful since the acquisition of relevant annotations/labels is very expensive in the medical field. Second, the model outputs description for a given CXR image instead of classifying it into a disease category. The consideration for this design is as follows. In clinical practice, it is not always feasible to draw solid conclusions on underlying diseases solely based on CXR images. For example, prominent/increased lung markings likely indicate an infection, chronic bronchitis, interstitial lung disease, heart failure, or normal aging. In case of such a symptom is observed, it is more appropriate to just describe the symptom, instead of giving a classification of diseases. The model is designed to follow this strategy. Moreover, this model's behavior

is similar to what radiologists usually do in their daily practice.

However, the model still requires improvement. Since the model is an end-to-end architecture that directly learns from reports and also generates reports, it does not explicitly give classification results. This makes it difficult to quantitatively evaluate the model performance. Currently we rely on human inspection for this purpose. We are dealing with this problem by adding a classification module in the neural network.

We believe the automated AI system developed in this work is useful and will greatly reduce the labor of doctors in the near future.

Acknowledgments

This work was funded by National Natural Science Foundation of China (Grant No. 11774158 to JZ, 11774157 to JW, 81720108022, 91649116 and 81571040 to BZ, and 11334004 to WW) ; by the social development project of science and technology project in Jiangsu Province (BE2016605 and BE201707 to BZ); in part by "Six big talent peak" high-level talent project (2016-WSN-160) from Jiangsu Provincial Department of Human Resources, Jiangsu Province, China; Key Project supported by Medical Science and technology development Foundation, Nanjing Department of Health(YKK18062,YKK17058) and the 66th batch of China Postdoctoral Science Foundation surface projects(2019M661805).The authors acknowledge HPCC of Nanjing University for the computational support.

Conflict of interests

The authors declare no conflict of interests.

References

- [1] Tudor G, Finlay D, and Taub N. An assessment of inter-observer agreement and accuracy when reporting plain radiographs. *Clinical radiology* 1997;52(3):235-238. doi: 10.1016/S0009-9260(97)80280-2
- [2] Rajpurkar P, Irvin J, Zhu K, Yang B, Mehta H, Duan T et al. Chexnet: Radiologist-level pneumonia detection on chest x-rays with deep learning. *arXiv preprint 2017;arXiv:1711.05225*.
- [3] Ypsilantis PP and Montana G. Learning what to look in chest x-rays with a recurrent visual attention model. *arXiv preprint 2017;arXiv:1701.06452*.
- [4] Pesce E, Ypsilantis PP, Withey S, Bakewell R, Goh V, and Montana G. Learning to detect chest radiographs containing lung nodules using visual attention networks. *arXiv preprint 2017;arXiv:1712.00996*.
- [5] Islam MT, Aowal MA, Minhaz AT, and Ashraf K. Abnormality detection and localization in chest x-rays using deep convolutional neural networks. *arXiv preprint 2017;arXiv:1705.09850*.
- [6] Yao L, Poblenz E, Dagunts D, Covington B, Bernard D, and Lyman K. Learning to diagnose from scratch by exploiting dependencies among labels. *arXiv preprint 2017;arXiv:1710.10501*.

- [7] Yan C, Yao J, Li R, Xu Z, and Huang J. Weakly supervised deep learning for thoracic disease classification and localization on chest x-rays. In Proceedings of the 2018 ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics 2018;103-110.
- [8] Guan Q, Huang Y, Zhong Z, Zheng Z, Zheng L, and Yang Y. Diagnose like a radiologist: Attention guided convolutional neural network for thorax disease classification. arXiv preprint 2018;arXiv:1801.09927.
- [9] Rubin J, Sanghavi D, Zhao C, Lee K, Qadir A and Xuwilson M. Large scale automated reading of frontal and lateral chest x-rays using dual convolutional neural networks. arXiv preprint 2018;arXiv:1804.07839.
- [10] Goodfellow I, Bengio Y, and Courville A. Deep learning (adaptive computation and machine learning series). Adaptive Computation and Machine Learning series 2016;800-800.
- [11] Silver D, Huang A, Maddison CJ, Guez A, Sifre L, Van DJ et al. Mastering the game of go with deep neural networks and tree search. Nature 2016;529:484-489. doi: 10.1038/nature16961
- [12] Xiong HY, Alipanahi B, Lee LJ, Bretschneider H, Merico D, Yuen RK et al. The human splicing code reveals new insights into the genetic determinants of disease[J]. Science 2015;347(6218):1254806-1254806. doi: 10.1126/science.1254806
- [13] Xu K, Ba J, Kiros R, Cho K, Courville A, Salakhudinov R et al. Show, attend and tell: Neural image caption generation with visual attention. In International conference on machine learning pages 2015;2048-2057.
- [14] Zhang Z, Xie Y, Xing F, McGough M, and Yang L. Mdnnet: A semantically and visually interpretable medical image diagnosis network. In Proceedings of the IEEE conference on computer vision and pattern recognition 2017;6428-6436. doi: 10.1109/CVPR.2017.378
- [15] Wang X, Peng Y, Lu L, Lu Z, and Summers RM. Tienet: Text-image embedding network for common thorax disease classification and reporting in chest x-rays. In Computer Vision and Pattern Recognition 2018;9049-9058.
- [16] Huang G, Liu Z, Maaten LVD, and Weinberger KQ. Densely connected convolutional networks. In Computer Vision and Pattern Recognition 2017;2261-2269. doi:10.1109/CVPR.2017.243
- [17] Hochreiter S and Schmidhuber J. Long short-term memory. Neural Computation 1997;9:1735-1780. doi:10.1162/neco.1997.9.8.1735
- [18] Zaremba W, Sutskever I, and Vinyals O. Recurrent neural network regularization. arXiv preprint 2014;arXiv:1409.2329.
- [19] Sun J. Chinese word segmentation module. <https://github.com/fxsjy/jieba>.
- [20] Wang X, Peng Y, Lu L, Lu Z, Bagheri M, and Summers RM. Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In Computer Vision and Pattern Recognition 2017;3462-3471. doi: 10.1109/CVPR.2017.369.
- [21] Vedantam R, Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. In Computer Vision and Pattern Recognition 2015;4566-4575. doi: 10.1109/CVPR.2015.7299087
- [22] Sutskever I, Vinyals O, and Le QV. Sequence to sequence learning with neural networks. In Advances

in neural information processing systems 2014;3104-3112.